arXiv:1703.06461v1 [math.OC] 19 Mar 2017

Regress-Later Monte Carlo for Optimal Inventory Control with applications in energy

Alessandro Balata* and Jan Palczewski†

School of Mathematics, University of Leeds, LS2 9JT, Leeds, United Kingdom

21st March 2017

Abstract

We develop a Monte-Carlo based numerical method for solving discrete-time stochastic optimal control problems with inventory. These are optimal control problems in which the control affects only a deterministically evolving inventory process on a compact state space while the random underlying process manifests itself through the objective functional. We propose a Regress Later modification of the traditional Regression Monte Carlo which allows to decouple inventory levels in two successive time steps and to include in the basis functions of the regression the dependence on the inventory levels. We develop a backward construction of trajectories for the inventory which enables us to use policy iteration of Longstaff-Schwartz type avoiding nested simulations.Our algorithm improves on the grid discretisation procedure largely used in literature and practice, and on the recently proposed control randomisation by Kharroubi et al. (2014). We validate our approach on two numerical examples: one is a benchmark problem of energy arbitrage used to compare different methods available in literature, the other is a high-dimensional problem of the management of a battery with the purpose of assisting the operations of a wind turbine in providing electricity to a group of buildings in a cost effective way.

1 Introduction

Our study of inventory control problems is motivated by the management of a limited amount of resources (which can be replenished or not) in order to maximise a measure of performance based on possibly both the inventory level itself and an exogenous process which for example could be the price or the demand for that particular resource. The inventory level influences the admissible controls since when the inventory is full it is not possible to accumulate resources further and analogously when it is empty it is not possible to withdraw them. Examples include management of energy storage such as batteries (Haessig et al., 2013), flywheels or pump-hydro stations for electricity (Zhaoa and Davison, 2009), underground caves for gas storage (Boogert and de Jong, 2008); or management of warehouse/shop stocks subject to fluctuations of product prices and demand (Secomandi, 2010).

In general, an inventory control problem involves two, possibly multi-dimensional, processes: an *underlying/exogeneous process* being a discrete-time Markov process (Markov chain) with an arbitrary state space and a controlled *inventory process* evolving deterministically on a compact state

*email: bn14a5b@leeds.ac.uk
†email: J.Palczewski@leeds.ac.uk

space. The set of admissible controls is usually state-dependent particularly due to the constraints on the inventory levels. The exogeneous process is not affected by controls but influences the set of admissible controls and the objective functional. In an energy arbitrage context where the storage is used to buy energy cheap and sell expensive, the underlying process is the energy price while the inventory process monitors the level of the storage. More complex problems are discussed in Section 5.

We intend to approach our problem in a probabilistic way, through its dynamical programming formulation and, in particular, we study algorithms which can be classified as regression Monte Carlo methods. Introduced in literature by Carriere (1996),Tsitsiklis and VanRoy (2001) and Longstaff and Schwartz (2001) these methods have proved their effectiveness in pricing of American and Bermudan options for complex dynamics of underlying instruments (Jain and Oosterlee, 2012). They borrow from standard Monte Carlo methods their independence on the dynamics as long as these can be simulated. This feature prompted interest in applying regression Monte Carlo to inventory control problems where the exogeneous process is commonly multi-dimensional and evolving on a non-compact state space preventing successful application of Markov chain approximation methods (Kushner and Dupuis, 2001). Although similarly to pricing of American/Bermudan options the control in inventory problems does not affect the random dynamics, the action of the control changes the future inventory level making the problem significantly more difficult. The simplest and most commonly used approach borrows from Markov chain approximations by discretising the inventory space and applying regression Monte Carlo independently at each point of the discrete grid, see, e.g., Boogert and de Jong (2008); Kitapbayev et al. (2015); Delage et al. (2017); Nadarajah et al. (2017). However, due to separate regressions needed at each discretisation point of the inventory, the method is inefficient and computational complexity becomes prohibitive when the dimension of the inventory space grows beyond two. Moreover, our numerical study shows that the computed strategy is significantly inferior to those obtained by other methods available in the literature and the ones proposed in this paper.

A natural improvement to the inventory space discretisation is to include the inventory variable in the regression in order to be able to read from the resulting regression surface all information "contained in each discretisation level" and, without interpolating, those in-between. Algorithms in Carmona and Ludkovski (2010); Denault et al. (2013) require that the set of available controls be finite and small (in their examples, they allow for up to three values of control: increase, decrease, hold). Particularly, the approach by Denault et al. (2013) is sensitive to the size of the control set as for each possible control value it runs a separate regression at each time step in the backward dynamic programming procedure. Kharroubi et al. (2014) treat the control as a state variable and approximate the continuation value in dynamic programming procedure as a function of three variables: the state of endogeneous process, the inventory and the control. This allows them to optimise over an infinite (or even continuous) set of controls at the cost of adding additional dimensions to the regression and basis functions. Langrené et al. (2015) shows an application of this methodology to natural resource extraction. In fact, Kharroubi et al. (2014) solve a more general problem of controlling diffusions, so that the control affects the random dynamics itself. We address this general problem in a companion paper Balata and Palczewski (2017).

This paper builds on a non-standard variation of regression Monte Carlo called Regress Later Monte Carlo (RLMC) which can be traced back to Broadie et al. (2000); Glasserman and Yu (2002); Broadie and Glasserman (2004) and was recently studied in Beutner et al. (2013); Jain and Oosterlee (2015); Nadarajah et al. (2017). In the classical regression Monte Carlo (Regress Now Monte Carlo, RNMC), the conditional expectation of future payoffs (continuation value) appearing in the dynamic programming equation is approximated by a combination of a finite number of basis functions of present state variables through minimisation of the least square error along simulated trajectories

of the underlying process. RLMC approximates first the value function at the following time by a finite combination of basis functions chosen in such a way that their conditional expectations can be computed analytically. Beutner et al. (2013) demonstrate that this modification of RNMC has a potential for significant improvement in the convergence rate of computed approximations to the value function. Nadarajah et al. (2017) compare RNMC and RLMC in the context of energy real options and find that the latter excels in valuation of the dual upper bound for the value function.

In the existing literature RLMC is predominantly regarded as a tool to reduce the approximation error in traditional exogenous Regression Monte Carlo problems (e.g., American option pricing). To the best of our knowledge, its ground-breaking potential for problems with endogeneous (controlled) state variables has not been recognised yet. Our *first contribution* is a formal derivation of the Regress Later Monte Carlo method for inventory problems avoiding discretisation of the inventory space as well as control randomisation; the reader is refered to a companion paper Balata and Palczewski (2017), where a general theory for controlled Markov chains is developed.

Policy iteration of Longstaff and Schwartz (2001) (here called *performance iteration* to distinguish it from classical policy iteration methods in control of Markov chains (Kushner and Dupuis, 2001)) improved upon Tsitsiklis and VanRoy (2001) by reducing propagation of the approximation error of the continuation value in the recursive formula for value function of an optimal stopping problem. In their algorithm, at each step of the dynamic programming backward procedure the continuation value is computed as a conditional expectation of realised payoffs obtained by following the estimate of the optimal strategy. This approach relies on the fact that the control (stopping) manifests itself through the objective functional but does not affect the dynamics of the underlying process. In inventory problems the level of inventory is determined by the control. Performance iteration would therefore require resimulation of inventory levels after an estimate of the optimal control is found in each backward step. The *second contribution* of this paper is the introduction of a technique for construction of optimal paths for the inventory backward in time, which considerably reduces the computations required by the performance iteration approach.

In the final part of the paper, we perform a numerical study of existing algorithms that use performance iteration for the solution of inventory control problems. For this we use a benchmark problem of energy arbitrage on a finite time horizon in which the electricity store is employed to buy energy when the energy price is low and to sell when the price is high. The exact strategy is intricate as it has to take into acccount the state of charge, the distribution of future prices, the time remaining for trading as well as the time-value of money (discounting). We demonstrate that the policy estimated by the algorithm involving discretisation of the inventory variable is significantly worse than those obtained either via control randomisation of Kharroubi et al. (2014) or via methods introduced in this paper. Although the differences between performance of the latter algorithms is much smaller, regress later Monte Carlo methods of this paper compute consistently better control policies than control randomisation.

We conclude the numerical section of the paper with a multi-dimensional problem motivated by Scott et al. (2014). It consists in controlling a battery within a system with a wind turbine, stochastic demand and bi-directional connection to the grid trading at stochastic exogeneous prices. The objective is to maximise the expected profits from operating the battery and involves using the battery to store electricity when the wind supply exceeds demand or when the grid price is low, whereas providing power to the demand or selling it to the grid when the price is high. The underlying stochastic process is three dimensional modelling wind power, demand and electricity price. The control is effectively two-dimensional with the inventory representing the state of charge of the battery. We demonstrate that regress-later Monte Carlo methods developed in this paper performs well in this multi-dimensional problem and exploit the statistical information about complicated dynamics of the

underlying processes in the computation of the control policy. We further compute the value of the battery showing the potential for practical applications in project valuation (c.f. Guthrie (2009); Swindle (2014)).

The paper is organised as follows. In Section 2 we introduce the inventory problem including a discussion of a link to a continuous-time model. In Section 3 we present the regression Monte Carlo-type approach with two alternative iterative methods: value and performance iteration. It is followed by an explanation why such methods cannot be directly applied to inventory problems and discussion of solution available in the literature. Section 4 contains our main contributions: a Regress Later Monte Carlo method for inventory problems and backward simulation of inventory process to avoid nested simulations within the context of performance iteration. In the last section we validate our approach on two numerical examples: one is a benchmark problem of energy arbitrage used to compare different methods available in literature, the other is a high-dimensional problem of the management of a battery with the purpose of assisting the operations of a wind turbine in providing electricity to a group of buildings in a cost effective way.

2 Inventory Control problem

This paper assumes *discrete-time dynamics* in the control problem. However, for presentation reasons we start with a continuous-time model and develop its approximation in discrete-time leading to a general discrete-time inventory control problem.

2.1 Motivating continuous-time problem

Consider a multidimensional process (X, I) with an exogenous (uncontrolled) p-dimensional random component $X_t(\omega) : \Omega \times [0,T] \to \mathbb{R}^p$ and an endogenous q-dimensional controlled component, $I_t(u) : \mathcal{U}^q \times [0,T] \to [0, I^1_{\max}] \times \dots [0, I^q_{\max}]$, where $\mathcal{U}^q$ is the family of admissible control processes to be specified in details below. We assume the dynamics

$$\begin{cases} dX_t = & \mu(t, X_t)\, dt + \Sigma(t, X_t)\, dW_t \\ dI_t = & \hat{\varphi}(t, u_t, I_t)\, dt\,, \quad u \in \mathcal{U}^q, \end{cases} \tag{1}$$

with coefficients μ and Σ such that a strong non-exploding solution exists for any initial value of X_0. Above, (W_t) is a multi-dimensional standard Brownian motion and Σ is a matrix of appropriate size. We further require that $\hat{\varphi}$ is a continuous function of its arguments. Denoting by $\mathcal{F}^X$ the filtration generated by the process (X_t), we specify the set of admissible controls $\mathcal{U}^q$ as all non-anticipating processess (u_t) (with respect to $\mathcal{F}^X$) satisfying $u_t = (u^1_t, \dots, u^q_t) \in [\underline{u}^1, \overline{u}^1] \times \dots \times [\underline{u}^q, \overline{u}^q]$ for all t and such that $I_t \in [0, I^1_{\max}] \times \dots \times [0, I^q_{\max}]$.

We introduce the following pathwise performance measure given the processes (X, I, u):

$$\hat{J}(t, X, I, u) = \int_t^T \hat{f}(s, X_s, I_s, u_s)\, ds + \hat{g}(X_T, I_T), \tag{2}$$

where $\hat{f}$ and $\hat{g}$ are measurable functions such that the expectation of $\hat{J}(t, X, I, u)$ is well-defined for any admissible control (u_t). Our aim is to maximise the expected value of the above measure over admissible control processes $(u_t) \in \mathcal{U}^q$:

$$V(t, x, i) = \sup_{u \in \mathcal{U}^q} \mathbb{E}\Big[\hat{J}(t, X, I, u)\Big| X_t = x, I_t = i\Big]. \tag{3}$$

Remark. Given the form of the functions $\hat{f}$ and $\hat{g}$ and the Markov property of the process (X, I) we can expect the optimal process u to be in a feedback form $u_t^*(X_t, I_t)$ for a Borel-measurable function u^*.

For simplicity of presentation, we fix $p = q = 1$, however all the developments of the paper can be easily extended to $p > 1$ and $q > 1$; only the backward construction of inventory path will be constrained to $q = 1$. We introduce now the discrete counterpart of the problem (1)-(3).

Consider a discretisation of the interval $[0, T]$ into N equispaced points such that $\Delta = T/N$, denote $\Delta W_n = W_{(n+1)\Delta} - W_{n\Delta}$. Using Euler approximation to the dynamics (1) we obtain a discrete-time controlled Markovian system:

$$\begin{cases} X_{n+1} = & X_n + \mu(n, X_n)\,\Delta + \Sigma(n, X_n)\,\Delta W_n \\ I_{n+1} = & I_n + \hat{\varphi}(n, u_n, I_n)\,\Delta\,, \quad u \in \mathcal{U}. \end{cases} \tag{4}$$

Here, $\mathcal{U}$ is a set of controls $(u_n)_{n=0,\dots,N-1}$ adapted to the filtration generated by (X_n) satisfying $u_n \in [\underline{u}, \overline{u}]$ and $I_n \in [0, I_{\max}]$ for all n. The value function

$$V(n, x, i) = \sup_{(u_s)\in\mathcal{U}} \mathbb{E}\Big[\sum_{s=n}^{N-1} \hat{f}(s, X_s, I_s, u_s)\Delta + \hat{g}(X_N, I_N) \Big| X_n = x, I_n = i \Big]$$

satisfies the following Bellman equation:

$$\begin{cases} V(n, x, i) = \sup\limits_{\substack{u\in[\underline{u},\overline{u}],\\ i+\hat{\varphi}(n,u,i)\Delta\in[0,I_{\max}]}} \Big\{ \hat{f}(n, x, i, u)\Delta + \mathbb{E}\big[V(n+1, X_{n+1}, \varphi(n, u, i))|X_n = x\big] \Big\}, \\ V(N, x, i) = \hat{g}(x, i). \end{cases} \tag{5}$$

2.2 General discrete-time model

In full generality, we assume that (X_n) is a Markov process on $\mathbb{R}^p$ and I is a controlled process with the dynamics specified as follows:

$$I_{n+1} = \varphi(n, u_n, I_n), \tag{6}$$

where φ is a Borel-measurable function. The above definition of the set of admissible controls $\mathcal{U}$ stays in force, i.e., a stochastic process $(u_n)_{n=0,\dots,N-1}$ is an admissible control if it is adapted to the filtration generated by (X_n), satisfies $u_n \in [\underline{u}, \overline{u}]$ and $I_n \in [0, I_{\max}]$ for all n.

In this setting, we define a performance measure

$$J(n, X, I, u) = \sum_{s=n}^{N-1} f(s, X_s, I_s, u_s) + g(X_N, I_N), \tag{7}$$

(note that we dropped Δ multiplying f as the system is stated in discrete-time only). The value function is given by $V(n, x, i) = \sup_{(u_s)\in\mathcal{U}} \mathbb{E}[J(n, X, I, u)|X_n = x, I_n = i]$. As the dynamics of (I_n) are deterministic and the control is adapted, we can exploit the measurability of I_{n+1} with respect to (X_n, I_n) in the Bellman equation for V:

$$\begin{cases} V(n, x, i) = \sup\limits_{\substack{u\in[\underline{u},\overline{u}],\\ \varphi(n,u,i)\in[0,I_{\max}]}} \Big\{ f(n, x, i, u) + \mathbb{E}\big[V\big(n+1, X_{n+1}, \varphi(n, u, i)\big)|X_n = x\big] \Big\}, \\ V(N, x, i) = g(x, i). \end{cases} \tag{8}$$

This property will be useful for numerical methods presented in Section 4.

Remark. Some applications could require to have randomness included in the dynamics of I; extension of this work to general controlled Markov chains is pursued in Balata and Palczewski (2017).

When designing a numerical algorithm for the solution of the dynamic problem (8) the most natural approach is that of computing $V(n, x, i)$ backward over time, starting from the known terminal condition. The problem, however, lies in the computation of the conditional expectation in (8) as the future state depends on the choice of control u. In the classical regression Monte Carlo, this would require redoing the regression step for each choice of the feedback function $u_n(X_n, I_n)$. In the following section, we discuss classical regression Monte Carlo in details and shed further light on the aforementioned problem with its application to an optimal control problem. We then sketch existing methods for overcoming this difficulty paving the way to our contribution.

3 Classical Regression MC for inventory optimisation

As explained in the introduction, Regression Monte Carlo (RMC) is one of the most flexible approaches to solve inventory control problems. It was introduced in Carriere (1996) and made famous by Longstaff and Schwartz (2001) in the context of American option pricing. The method is very powerful for three main reasons:

- the method is applicable for any Markov process X as long as it can be simulated;
- the complexity, excluding the regression step, scales linearly with the number of dimensions;
- the regression step allows to compute the conditional expectations without having to run a nested MC simulations for each state.

Remark. The first two points come from the use of Monte Carlo simulations, while the last one is the result of the introduction of regression instead of resimulation. Note that the regression step can potentially scale exponentially with the number of dimensions, but significantly slower than the grid discretisation methods.

This is made possible because the RMC method approximates conditional expectations using a linear combination of basis functions of state variables. The choice of basis functions plays a crucial role in the accuracy and computational complexity of approximating conditional expectations. When multi-dimensional basis functions are constructed by multiplying basis functions for individual dimensions, the complexity of the regression step grows exponentially with the number of directions. This can, however, be overcome with a careful choice of multi-dimensional basis functions based on domain knowledge or preliminary numerical results.

3.1 Regression approximation of conditional expectation

Recall that in the formal presentation of the method for inventory problems, we restrict attention to the one-dimensional case: a single inventory and a one-dimensional random state process (X_n). Notice that the higher dimensional case is identical; in fact the dimension of X and I does not influence any of the arguments that follow while only in Section 4.3 we require the dimension of I to be one. Define $\mathcal{L}_\phi^K$ as the linear space generated by the set of (linearly independent) functions $\{\phi_k(\cdot,\cdot)\}_{k=1}^K$: $\mathbb{R} \times [0, I_{\text{max}}] \to \mathbb{R}$. The dynamic programming iteration (8) can be computed in the regression MC framework taking advantage of the cross sectional information comprised in the different paths and projecting the value function at time $n+1$ over $\mathcal{L}_\phi^K(n)$, where with this notation we mean the set

of basis function evaluated at time n, i.e. $(\phi_k(X_n, I_n))_{k=1,\dots,K}$. More generally, for any functional $h : (\mathbb{R} \times [0, I_{\max}])^{N-n} \to \mathbb{R}$ we can compute an approximation of the conditional expectation $\hat{\mathbb{E}}_{n,x,i}\big[h((X_t, I_t)_{t=n+1,\dots,N})\big] \approx \mathbb{E}\big(h((X_t, I_t)_{t=n+1,\dots,N})|X_n = x, I_n = i\big)$ as follows:

$$\hat{\mathbb{E}}_{n,x,i}\Big[h\big((X_t, I_t)_{t=n+1,\dots,N}\big)\Big] = \sum_{k=1}^{K} \alpha_k \phi_k(x, i), \tag{9}$$

where coefficients (α_k) solve the optimisation problem

$$\mathbb{E}\bigg[\mathbb{E}\Big[h\big((X_t, I_t)_{t=n+1,\dots,N}\big)\Big|X_n, I_n\Big] - \sum_{k=1}^{K} \alpha_k \phi_k(X_n, I_n)\bigg]^2 \to \min. \tag{10}$$

In the above, the outer expectation averages over the distribution of (X_n, I_n) and we intentionally supress in the notation the dependence of I_{n+1} on the control u_n for the reasons that will become clear while we describe existing methods for dealing with the limitations of the regression MC for controlled processes. We only hint here that one of these methods involves introducing randomness to the dynamics of (I_n).

A practical implementation of the RMC method requires generation of a set of M trajectories $(Z_n^m, Y_n^m)_{n,m=1}^{N,M}$ of the process $(X_n, I_n)_{n=1}^{N}$ to approximate expectations in (10). Notice that, because the inventory is controlled, the generation of the Y's is problematic, but assume for now that it is possible. Coefficients $(\alpha_k^n)_{k=1,\dots,K}$ in the approximation (9) of the conditional expectation $\mathbb{E}\big(h((X_t, I_t)_{t=n+1,\dots,N})|X_n = x, I_n = i\big)$ are then computed as

$$\alpha^n = \underset{\alpha\in\mathbb{R}^K}{\operatorname{argmin}}\Big\{\sum_{m=1}^{M}\Big[h((Z_t^m, Y_t^m)_{t=n+1,\dots,N}) - \sum_{k=1}^{K}\alpha_k\phi_k(Z_n^m, Y_n^m)\Big]^2\Big\}. \tag{11}$$

This is a standard linear regression procedure where M responses $\big\{h((Z_t^m, Y_t^m)_{t=n+1,\dots,N})\big\}_{m=1}^{M}$ are regressed against the set of K explanatory variables $\big\{\phi_k(Z_n^m, Y_n^m)\big\}_{m=1}^{M}$ indexed by $k = 1, \dots, K$.

3.2 Value and Performance iteration

Value iteration for optimal stopping problems has been introduced in the literature on regression Monte Carlo by Tsitsiklis and VanRoy (2001). In the optimal control case, the update rule for the value function follows directly from equation (8):

$$V(n, x, i) = \sup_{\substack{u\in[\underline{u},\overline{u}],\\ \varphi(n,u,i)\in[0,I_{\max}]}} \Big\{f(n, x, i, u) + \hat{\mathbb{E}}_{n,x,i,u}\big[V(n + 1, X_{n+1}, I_{n+1})\big]\Big\}, \tag{12}$$

where the estimator $\hat{\mathbb{E}}$ introduced in (9) has been extended to accomodate the control u at time n; recall that this control acts in a deterministic way on I_n only. The difficulties arising in extending Tsitsiklis and VanRoy (2001) approach to inventory problems are discussed in Section 3.3. In a nutshell, the future value of I_{n+1} depends on the control applied at time n therefore preventing the computation of the conditional expectation estimator $\hat{\mathbb{E}}_{n,x,I_n}\big[V(\cdot, I_{n+1})\big]$.

Policy iteration for regression Monte Carlo was introduced in Longstaff and Schwartz (2001) for optimal stopping problems; despite its name, this approach is different to the well known policy iteration algorithm for the control of Markov chains (Kushner and Dupuis, 2001) in which a parametrised

control is improved iteratively. To avoid misunderstanding, throughout the paper we will refer to the Longstaff-Schwartz policy iteration approach as *performance iteration*. Indeed, in this approach the update rule for the value function is based on the dynamic programming equation for the performance measure, instead of that for the value function:

$$V(n,x,i) = \mathbb{E}\Big[f(n,x,i,u_n^*) + J\big(n+1,(X_t,I_t,u_t^*)_{t=n+1,\dots,N}\big)|X_n = x, I_{n+1} = \varphi(n,u,i)\Big], \quad (13)$$

where the "realised profits" (performance measure) are given by a recurrence relation

$$J(n,(X_t,I_t,u_t^*)_{t=n,\dots,N}) = f(X_n,I_n,u_n^*(X_n,I_n)) + J(n+1,(X_t,I_t,u_t^*)_{t=n+1,\dots,N}) \quad (14)$$

with $J(N,(X_N,I_N,u_N^*)) = g(X_N,I_N)$. For the simplicity of notation, in parameters of J we include u_N^* which is never used since there is no control excerted at the terminal time N; for the sake of clarity, we will put $u_N^*(x,i) = 0$. An optimal control u_n^* is estimated via

$$u_n^*(x,i) = \underset{\substack{u\in[\underline{u},\overline{u}],\\ \varphi(n,u,i)\in[0,I_{\max}]}}{\operatorname{argmax}} \Big\{f(n,x,i,u) + \hat{\mathbb{E}}_{n,x,i,u}\big[J(n+1,(X_t,I_t,u_t^*)_{t=n+1,\dots,N})\big]\Big\}. \quad (15)$$

Note that a direct implementation of equation (13) requires computation of

$$J(n+1,(X_t,I_t,u_t^*)_{t=n+1,\dots,N}) = \sum_{t=n+1}^{N-1} f(k,X_t,I_t,u_t^*) + g(X_N,I_N)$$

for each choice of u_n^*, which numerically means resimulating the path from time n to the terminal time incurring heavy computations. Longstaff and Schwartz (2001) in an optimal stopping context propose a more efficient update rule based on (14) in the sense that, given a path of (X,I), it is possible to directly use the value of $J(n+1,(X_t,I_t,u_t^*)_{t=n+1,\dots,N})$ computed at the previous time step to update $J(n,(X_t,I_t,u_t^*)_{t=n,\dots,N})$. In the control case, this however requires that the value of $J(n+1,\cdot)$ is known for the particular trajectory $(X_t,I_t,u_t^*)_{t=n+1,\dots,N}$ which itself depends on the control u_n^* applied at time n. This seemingly invalidates the improvement of Longstaff and Schwartz (2001) that avoids resimulation. We show in Section 4 how this can be overcome.

In many problems, given a good approximation of the conditional expectations, both value and performance iteration can achieve precise estimates. Sometimes however, either when the value function has a "difficult" shape or because the running profit per period is relatively small compared to the continuation value, it is better to use performance iteration in order to minimise errors and improve the quality of the estimated policy. This approach works very well because the estimates of the conditional expectations are used only to take decisions while the value of these decisions is computed exactly, forward in time. Despite being time consuming the performance iteration can sometimes be the only possible alternative, in particular when using polynomial basis function because, even though they could be good enough to take decisions, they may be insufficient for accurate representation of the value of those. It should be pointed out however that the empirical regression calculated in the performance iteration case can be unstable, suffering from the relatively high variance of the training points. In the value iteration case, on the other hand, response points are typically close to the previously estimated value function, causing the propagation of the regression error over time and, therefore, leading to stable, but biased, regression coefficients.

Finally note that both value and performance iteration are designed to compute an estimation of the optimal policy (u_t^*) and, in fact the value function obtained as output of the backward iterations

should not be used as an approximation of the performance of this policy. Rather, we run an additional forward Monte Carlo simulation to establish the value of the estimated policy. The only output of the backward procedure should therefore be the matrix of regression coefficients $A = \{\alpha_k^n\}_{n,k=1}^{N,K}$, to be used in forward simulations, to take decisions. The value of the policy induced by A is computed as

$$v(x,i) = \frac{1}{M'} \sum_{m=1}^{M'} \Big(\sum_{n=1}^{N} f(n, Z_n^m, Y_n^m, \nu_n^m) + g(Z_n^m, Y_N^m) \Big), \tag{16}$$

where the inventory process evolves as $Y_{n+1}^m = \varphi(Z_n^m, Y_n^m, \nu_n^m)$ with $Z_1^m = x$, $Y_1^m = i$ and

$$\nu_n^m = \underset{\substack{u \in [\underline{u}, \overline{u}], \\ \varphi(n,u,Y_n^m) \in [0, I_{\max}]}}{\operatorname{argmax}} \Big\{ f(n, Z_n^m, Y_n^m, u) + \sum_{k=1}^{K} \alpha_k^n \phi_k(Z_n^m, Y_n^m) \Big\}. \tag{17}$$

Note that since the process (X) is uncontrolled the whole set of trajectories $(Z_n^m)_{n,m=1}^{N,M'}$ in equation (16) can be precomputed. In the following, when presenting the algorithms, we use the notation "Evaluate the policy" to refer to the routine specified by equations (16)-(17).

Remark. It can be easily shown that, because the estimated policy must be not better than optimal, it is true that:

$$\mathbb{E}[v(x,i)] \leq V(1,x,i)$$

3.3 Limitations of classical RMC for inventory problems

In practical implementations of value or performance iteration, the approximation of conditional expectation operator $\hat{E}_{n,x,i}(\cdot)$ is computed using Monte Carlo approach as in (11). In the value iteration case, see (12), this amounts to regress a set of response points $\mathcal{Y} = \big(V(n+1, Z_{n+1}^m, Y_{n+1}^m)\big)_{m=1}^{M}$ over the sets of observations $\mathcal{B} = \big(\mathcal{X}_1, \ldots, \mathcal{X}_K\big)$, where $\mathcal{X}_k = \big(\phi_k(Z_n^m, Y_n^m)\big)_{m=1}^{M}$ for $k = 1, \ldots, K$, and $\big(Z_t^m, Y_t^m\big)_{m,t=1}^{M,N}$ is the set of simulated trajectories of $(X_t, I_t)_{t=1}^{N}$. However, contrary to the uncontrolled case studied by Tsitsiklis and VanRoy (2001) and Longstaff and Schwartz (2001), the process I depends on the control, so one cannot compute first the approximation of $\hat{\mathbb{E}}_{n,x,i}\big[V(n+1, X_{n+1}, I_{n+1})\big]$ and then optimise over controls. For each choice of control, one gets a different value of I_{n+1}. This yields two problems:

a) The calculation of an optimal control at time n is in itself a fixed point problem: one has to find such a control u_n^* that after doing a regression approximation of the conditional expectation with I_{n+1} determined by this control, the maximiser in (12) equals to u_n^*;

b) The value of $V(n+1, X_{n+1}, I_{n+1})$ has to be established for multiple outcomes of I_{n+1} achievable from I_n when applying an admissible control u_n. This prevents usage of previously computed values of $V(n+1, X_{n+1}, \cdot)$.

In the following subsections we briefly review methods available in literature to address the above issues. Our contribution is presented in Section 4.

3.4 Grid discretisation of inventory levels

Consider a discretisation $\Lambda_L = \{\lambda_1 = 0, \ldots, \lambda_L = I_{\max}\}$ of the state space of the controlled inventory, commonly an uniform grid $\Lambda_L = \{(l-1)\frac{I_{\max}}{L-1} : l = 1, \ldots, L\}$. Simulate M trajectories

Algorithm 1 Regression Monte Carlo algorithm with Grid Discretisation (GD)

Generate training points $\{Z_n^m\}_{n=1,m=1}^{N,M}$
Initialise the value function $V(N, Z_N^m, \lambda_l) = g(Z_N^m, \lambda_l), \quad j = 1, \dots, L,\ m = 1, \dots, M$
for $n = N-1$ to 1 **do**
    for $l = 1$ to L **do**
        $\alpha_n^l = \underset{\alpha \in \mathbb{R}^K}{\operatorname{argmin}} \{\sum_{m=1}^M [V(n+1, Z_{n+1}^m, \lambda_l) - \sum_{k=1}^K \alpha_k \phi_k(Z_n^m)]^2\}$
        $\hat{\mathcal{E}}(Z_n^m, \lambda_l) = \sum_{k=1}^K \alpha_n^{k,l} \phi_k(Z_n^m), \quad \forall m$
    Extend $\hat{\mathcal{E}}(Z_n^m, i)$ to $i \in [0, I_{\max}]$ by interpolation
    $u_{n,m}^l = \underset{\substack{u \in [\underline{u}, \overline{u}], \\ \varphi(n, u_n, \lambda_l) \in [0, I_{\max}]}}{\operatorname{argmax}} \{f(n, Z_n^m, \lambda_l, u) + \hat{\mathcal{E}}(Z_n^m, \varphi(n, u, \lambda_l))\}, \quad \forall m, l$
    if value iteration **then**
        $V(n, Z_n^m, \lambda_l) = f(n, Z_n^m, \lambda_l, u_{n,m}^l) + \hat{\mathcal{E}}(Z_n^m, \varphi(n, u_{n,m}^l, \lambda_l))\}, \quad \forall m, l$
    if performance iteration **then**
        extend $V(n+1, Z_{n+1}^m, \lambda^l)$ to $\underline{V}(n+1, Z_{n+1}^m, i), i \in [0, I_{\max}]$, by interpolation
        $V(n, Z_n^m, \lambda_l) = f(n, Z_n^m, \lambda_l, u_{n,m}^l) + \underline{V}(n+1, Z_{n+1}^m, \varphi(n, u_{n,m}^l, \lambda_l)), \quad \forall m, l$
Evaluate the policy

$(Z_t^m)_{m,t=1}^{M,N}$ of the uncontrolled process $(X_t)_{t=1}^N$. Note that the use of the grid is effectively equivalent to having $\{Y_n^{m,l} = \lambda_l\}_{m,l=1}^{M,L}$. In order to compute the continuation value in (12) regress the set of responses $\{V(n+1, Z_{n+1}^m, \lambda_l)\}_{m=1}^M$ over $\{\phi_k(Z_n^m)\}_{m,k=1}^{M,K}$ for each $l = 1, \dots, L$ obtaining a collection of coefficients $\{\alpha_n^{k,l}\}_{k,l=1}^{K,L}$ and approximations $\sum_{k=1}^K \alpha_n^{k,l} \phi_k(x)$ of $\mathbb{E}[V(n+1, X_{n+1}, \lambda_l) | X_n = x]$. Since these approximations are for a fixed set of inventory levels, an interpolation between those levels is needed in order to run the backward procedure in dynamic programming equation (12). Denoting the resulting approximation of the conditional expectation $\mathbb{E}[V(n+1, X_{n+1}, i) | X_n = x]$ by $\mathcal{E}(n, x, i)$, we obtain an approximate backward dynamic programming equation

$$V(n, Z_n^m, \lambda_l) = \sup_{\substack{u \in [\underline{u}, \overline{u}], \\ \varphi(n, u_n, \lambda_l) \in [0, I_{\max}]}} \Big\{ f(n, Z_n^m, \lambda_l, u) + \mathcal{E}(n, Z_n^m, \varphi(n, u, \lambda_l)) \big] \Big\}$$

for $l = 1, \dots, L$. Algorithm 1 provides details of the implementation. Extension to performance iteration results in a minor amendment as can be seen in the algorithm.

The inventory discretization method is wide-spread because of the relative simplicity of extending a traditional Regression Monte Carlo algorithm to controlled dynamics, see, e.g., Boogert and de Jong (2008); Delage et al. (2017); Kitapbayev et al. (2015); Nadarajah et al. (2017). However, due to separate regressions needed at each discretisation point of the endogenous (controlled) component, it could be prohibitive to implement it if there are more than two endogenous dimensions. As our empirical analysis shows, it is also inefficient even in a one-dimensional case.

A natural improvement to the discretisation technique would be that of including the inventory in the regression in order to be able to read from the resulting regression surface all information "contained in each discretisation level" and, without interpolating, those in-between. However, as discussed in Section 3.3, this is not easily achievable and only recently some alternatives have appeared in the literature, see Carmona and Ludkovski (2010); Denault et al. (2013); Kharroubi et al. (2014). Of the aforementioned papers only Kharroubi et al. (2014) provides a justification of the proposed algorithm. While it focuses on the controlled diffusion problem, Langrené et al. (2015) shows an

application of this methodology to natural resource extraction (i.e., with $\bar{u} = 0$ in our notation).

In the following subsections we discuss methodological developments of the papers mentioned in the previous paragraph with the main focus on how the authors solve difficulties highlighted in Section 3.3.

3.5 Quasi-simulation of the inventory

In order to include the inventory in the regression, its trajectories (I_n) have to be simulated before the control is computed. However, as the control directly guides the evolution of (I_n), the inventory should rather be stated at time N and then its paths constructed backward in time as the dynamic programming procedure progresses. Carmona and Ludkovski (2010) and Denault et al. (2013) resolved this issue in similar ways in the framework of an optimal switching problem related to the management of an energy storage device for which only 3 controls are available: charge, discharge and store; i.e. $u \in \{-1, 0, 1\}$. The small and finite size of the set of possible control is paramount for these methods.

In Carmona and Ludkovski (2010) the authors deal for the first time with the task of regressing $\{V(n+1, Z^m_{n+1}, Y^m_{n+1})\}_{m=1}^M$ over both exogenous and inventory processes. Clearly only three predecessors Y^m_n are possible and the authors suggest to test separately whether Y^m_{n+1} is an optimal successor of one of these three. This is achieved by regressing the value function at time $n+1$ over observations $(Z^m_n, Y^m_{n+1})_{m=1}^M$ and later testing whether an optimal control at time n for any of the three predecessors of Y^m_{n+1} drives the inventory process to Y^m_{n+1}. If there is such a predecessor in the interval $[0, I_{\max}]$, the trajectory Y^m is extended to time n and the performance iteration is used via a direct update as explained in Section 3.2. Otherwise, the Y^m_n is randomly placed in $[0, I_{\max}]$ and a value iteration approximation is used.

The approach in Denault et al. (2013) to construct backward paths of the inventory process uses basis functions dependent on X_n and I_n and requires the computation of 3 separate regressions for each choice of control at step n. This procedure does not guarantee that the trajectories $(Y^m_n)^{N,M}_{n,m}$ will remain inside $[0, I_{\max}]$ and they decided therefore to introduce an artificial penalty term and a random recollocation of the point Y^m_n once it is too far from the interval $[0, I_{\max}]$.

Reliable and thorough tests are not available, however, our experience and the literature seem to suggest that the grid discretisation method is preferred over quasi-simulation whenever the dimensionality of the problem is small enough to allow its use.

3.6 Control Randomisation

Control randomisation was introduced in Kharroubi et al. (2014); it differs from the previous methods in that the control becomes a state variable and it is simulated (according to some arbitrary model that ensures admissibility) along trajectories of $(X_t, I_t)_{t=1}^N$. Data underlying the regression Monte Carlo consists of three sets of trajectories $\{Z^m_n, Y^m_n, \tilde{u}^m_n\}^{N,M}_{n,m=1}$. In the case of value iteration (12), $\{V(n+1, Z^m_{n+1}, Y^m_{n+1})\}_{m=1}^M$ is regressed against basis functions evaluated at the points $\{Z^m_n, Y^m_n, \tilde{u}^m_n\}_{m=1}^M$. These regression basis functions are dependent now on the random control $\tilde{u}_n$, in addition to X_n and I_n so that the estimated continuation value will depend on the choice of the control (which is different on each sample trajectory). An optimal control at time n given $X_n = Z^m_n$ and $I_n = Y^m_n$ is approximated by the expression (if the maximum is attained)

$$u^*_n(x,i) = \underset{\substack{u\in[\underline{u},\overline{u}],\\ \varphi(n,u,i)\in[0,I_{\max}]}}{\operatorname{argmax}} \Big\{ f(n,x,i,u) + \hat{\mathbb{E}}_{n,x,i,u}\big[V(n+1, X_{n+1}, I_{n+1}))\big] \Big\}, \tag{18}$$

Algorithm 2 Regression Monte Carlo algorithm: Control Randomisation (CR)

1: Generate a random control $\{\tilde{u}_n^m\}_{n=1,m=1}^{N,M} \in \mathcal{U}$ along the training points $\{Z_n^m, Y_n^m\}_{n=1,m=1}^{N,M}$.
2: Initialise the value function $V(N, Z_N^m, Y_N^m) = g(Z_N^m, Y_N^m), \quad \forall m$
3: **for** n=N-1 to 1 **do**
4: $\alpha_{n+1} = \underset{\alpha \in \mathbb{R}^K}{\operatorname{argmin}} \{\sum_{m=1}^M [V(n+1, Z_{n+1}^m, Y_{n+1}^m) - \sum_{k=1}^K \alpha_{n+1}^k \phi_k(Z_n^m, Y_n^m, \tilde{u}_{n,m})]^2\}$
5: **if** value iteration **then**
6: For all m do

$$V(n, Z_n^m, Y_n^m) \tag{19}$$
$$= \sup_{\substack{u \in [\underline{u}, \overline{u}], \\ \varphi(n,u,Y_n^m) \in [0, I_{\max}]}} \left\{ f(n, Z_n^m, Y_n^m, u) + \sum_{k=1}^K \alpha_{n+1}^k \phi_k(Z_n^m, Y_n^m, u) \right\}$$

7: **if** performance iteration **then**
8: For each m, update trajectories $\{Y_j^m\}_{j=n+1}^N$ using control given by Equation (18)
9: $V(n, Z_n^m, Y_n^m) = \sum_{j=n}^{N-1} f\left(j, Z_j^m, Y_j^m, u_n^*\left(Z_j^m, Y_j^m\right)\right) + g(Z_N^m, Y_N^m)$
10: **Evaluate the policy**

where, with a slight abuse of notation, we included in the approximate conditional expectation $\hat{\mathbb{E}}$ the dependence on the control u. In general, multiple runs of the method could be needed to obtain precise estimates because the initial choice of the dummy control could drive the training points far from where the optimal control would have directed them.

Remark. In many practical applications, the mapping $u \mapsto \varphi(x, i, u)$ is injective (strictly increasing), hence there is a one-to-one correspondence between u_n and I_{n+1} given I_n. The addition of this new coordinate in the regression may be replaced by regressing $\{V(n+1, Z_{n+1}^m, Y_{n+1}^m)\}_{m=1}^M$ against basis functions evaluated at points $\{Z_n^m, Y_n^m, Y_{n+1}^m\}_{m=1}^M$. This shows that trajectories $\{Y_n^m\}_{n,m=1}^{N,M}$ may be constructed deterministically to offer maximum coverage of those parts of the state space in which the precision is the most desirable.

For further reference see Langrené et al. (2015) where the authors apply the numerical method in the context of natural resource extraction. For details about the implementation and extension to the performance iteration case refer to Algorithm 2. Notice that a resimulation stage, lines 9 and 10, is required in the performance iteration algorithm.

4 Regress Later Monte Carlo: a decoupling approach

In the previous section the regression was used to approximate the conditional expectation with respect to (X_n, I_n) directly as a linear combination of basis functions of those two variables. As opposed to this *regress-now* approach, we employ in this section a *regress-later* idea in which the conditional expectation with respect to (X_n, I_n) is computed in two stages. First, a conditional expectation with respect to (X_{n+1}, I_{n+1}) is approximated in a regression step by a linear combination of basis functions of (X_{n+1}, I_{n+1}). Then, analytical formulas (or precomputed numerical approximations) are applied to condition this linear combination of functions of future values on present values (X_n, I_n).

4.1 Regress later approximation

Unlike the regress-now method for approximating conditional expectations, the regress-later approach imposes conditions on basis functions:

Assumption 1. *For each basis function ϕ_k, $k = 1, \dots, K$, the conditional expectation*

$$\hat{\phi}_k^n(x, i) = \mathbb{E}[\phi_k(X_{n+1}, i)|X_n = x]$$

can be computed analytically.

Remark.

1. Assumption 1 can be relaxed in practical numerical implementations because the conditional expectations of the basis functions can be numerically precomputed.
2. In many problems, the process (X_n) is time homogeneous and, therefore, functions $\hat{\phi}_k^n$ do not depend on n. For simplicity of notation we drop the superscript in the following.
3. Multidimensional basis functions are often constructed as products of one-dimensional functions: $\phi_k(x, i) = \phi_k^X(x)\phi_k^I(i)$. Then the function $\hat{\phi}_k$ factors as follows:

$$\hat{\phi}_k(x, i) = \phi_k^I(i)\, \mathbb{E}[\phi_k^X(X_{n+1})|X_n = x].$$

We will now present regress-later solution to *value iteration* procedure (12). Assume that at time $n+1$ the value function $V(n+1, \cdot)$ has been computed for a set of training points $\{Z_{n+1}^m, Y_{n+1}^m\}_{m=1}^M$. We perform a regression to approximate $V(n+1, \cdot)$ with a linear combination of basis functions:

$$V(n+1, x, i) \approx \sum_{k=1}^{K} \alpha_k^{n+1} \phi_k(x, i),$$

where

$$\alpha^{n+1} = \underset{\alpha \in \mathbb{R}^K}{\operatorname{argmin}} \left\{ \sum_{m=1}^{M} \Big[V(n+1, Z_{n+1}^m, Y_{n+1}^m) - \sum_{k=1}^{K} \alpha_k \phi_k(Z_{n+1}^m, Y_{n+1}^m) \Big]^2 \right\}. \tag{20}$$

Moving one step backward to time n, we would like to compute the value function $V(n, \cdot, \cdot)$ and an optimal control $u_n(\cdot, \cdot)$. We select a set of training points $\{Z_n^m, Y_n^m\}_{m=1}^M$, which can be placed independently from $\{Z_{n+1}^m, Y_{n+1}^m\}_{m=1}^M$, therefore, removing the limitation of the regress-now approach for which points at different times must be linked through their true dynamics. Here, these training points are only needed for further computation of an approximation of the value function at time n with a linear combination of basis functions. An approximation of an optimal control is computed as (assuming that the maximum is attained)

$$u_n^*(x, i) = \underset{\substack{u \in [\underline{u}, \overline{u}], \\ \varphi(n,u,i) \in [0, I_{\max}]}}{\operatorname{argmax}} \left\{ f(n, x, i, u) + \sum_{k=1}^{K} \alpha_k^{n+1} \hat{\phi}_k^n\big(x, \varphi(n, u, i)\big) \right\}$$

with the value function at sampling points estimated by

$$V(n, Z_n^m, Y_n^m) = f(n, Z_n^m, Y_n^m, u_n^m) + \sum_{k=1}^{K} \alpha_k^{n+1} \hat{\phi}_k^n\big(Z_n^m, \varphi(n, u_n^m, Y_n^m)\big)$$

with $u_n^m = u_n^*(Z_n^m, Y_n^m)$.

The above procedure can be easily adapted to *performance iteration* approach, (13), with resimulation. Assume that a control mapping $u_j^* : \mathbb{R} \times [0, I_{\max}] \to \mathbb{R}$ is known for $j = n+1, \dots, N-1$. Assume further that $(Y_j^m)_{j=n+1,\dots,N}$ correspond to those controls, i.e., $Y_{j+1}^m = \varphi\big(j, u_j^*(Z_j^m, Y_j^m), Y_j^m\big)$ for $j = n+1, \dots, N-1$. We approximate the conditional expectation $\mathbb{E}\big[J(n+1, (X_j, I_j, u_j^*)_{j=n+1,\dots,N}) \big| X_{n+1} = x, I_{n+1} = i\big]$ by a linear combination of basis functions $\sum_{k=1}^K \alpha_k^{n+1} \phi_k(x, i)$ performing a least-squares regression:

$$\alpha^{n+1} = \operatorname*{argmin}_{\alpha \in \mathbb{R}^K} \Big\{ \sum_{m=1}^M \Big[J\big(n+1, (Z_j^m, Y_j^m, u_j^*)_{j=n+1,\dots,N}\big) - \sum_{k=1}^K \alpha_k \phi_k(Z_{n+1}^m, Y_{n+1}^m) \Big]^2 \Big\}. \quad (21)$$

Note that although $J(n+1, \cdot)$ is not itself measurable with respect to (X_{n+1}, I_{n+1}), as is the case for the value iteration algorithm, the least-squares regression in (21) approximates the aforementioned conditional expectation. Indeed, the above regression projects $J(n+1, \cdot)$ on the linear space spanned by functions $\phi_k(X_{n+1}, I_{n+1})$, $k = 1, \dots, K$. This can be viewed as a superposition of two projections: the first one on the space of generated by functions of (X_{n+1}, I_{n+1}) and further on its linear subspace generated by basis functions $\phi_1, \dots, \phi_K$.

A backward step of the performance iteration algorithm computes a control mapping u_n^*:

$$u_n^*(x, i) = \operatorname*{argmax}_{\substack{u \in [\underline{u}, \overline{u}], \\ \varphi(n,u,i) \in [0, I_{\max}]}} \Big\{ f(n, x, i, u) + \sum_{k=1}^K \alpha_k^{n+1} \hat{\phi}_k^n\big(x, \varphi(n, u, i)\big) \Big\}. \quad (22)$$

For all points (Z_n^m, Y_n^m) the trajectories $\{Y_j^m\}_{j=n+1}^N$ must be recomputed as it is unlikely that the inventory process governed by $u_n^*(Z_n^m, Y_n^m)$ matches exactly Y_{n+1}^m. In Section 4.3 we show how to position training points for the inventory at time n to avoid this resimulation step.

4.2 Algorithm

Algorithm 3 lists a pseudocode for value and performance iteration using regress-later approach. Notice that the absence of interdependence among the training points $\{Z_n^m, Y_n^m\}_{n=1}^N$ at different times allows, at each time step, to place the training data where it is more convenient or needed. However, the optimal design of the training set is beyond the scope of the present paper and we generate observations Y_n^m of I_n from a uniform distribution on $[0, I_{\max}]$, while X_n^m are obtained from the simulation of the uncontrolled process X.

We would like to mention a method for the flexible design of basis functions for which the conditional expectation $\hat{\phi}_k^n$ can be computed analytically. The method has been presented first in Bouchard and Warin (2012) in the context of option pricing. It consists of dividing the state space into hypercubes and choosing affine functions within them. This approach allows to capture complex shapes as long as the number of hypercubes is sufficiently large. In our framework, this corresponds to choosing basis functions $\phi_k(x, i) = (a_k x + b_k i + c_k) 1_{\{(x,i) \in B_k\}}$, with disjoint hypercubes B_k.

4.3 Backward construction of inventory levels for performance iteration

The main bottleneck in using performance iteration for inventory problems is the need to resimulate the inventory level until maturity after each backward step. When the algorithm moves closer to the initial time, the length of trajectories that need to be resimulated grows. Eventually, the computational

Algorithm 3 Regress-later Monte Carlo algorithm (RLMC)

Generate training points $\{Z_n^m\}_{n=1,m=1}^{N,M}$ accordingly to the dynamics induced by X
Choose training points $\{Y_n^m\}_{n=1,m=1}^{N,M}$
Initialise the value function $V(N, Z_N^m, Y_N^m) = g(Z_N^m, Y_N^m), \quad \forall m$
for n=N-1 to 1 **do**
    $\alpha_{n+1} = \underset{\alpha\in\mathbb{R}^K}{\operatorname{argmin}} \left\{ \sum_{m=1}^M [V(n+1, Z_{n+1}^m, Y_{n+1}^m) - \sum_{k=1}^K \alpha_k \phi_k(Z_{n+1}^m, Y_{n+1}^m)]^2 \right\}$
    if value iteration **then**
        For all m do

$$V(n, Z_n^m, Y_n^m) \tag{23}$$
$$= \sup_{\substack{u\in[\underline{u},\overline{u}],\\ \varphi(j,u,Y_j^m)\in[0,I_{\max}]}} \left\{ f(n, Z_n^m, Y_n^m, u) + \sum_{k=1}^K \alpha_{n+1}^k \hat{\phi}_k^n(Z_n^m, \varphi(n, u, Y_n^m)) \right\}$$

    if performance iteration **then**
        Define control mapping u_n^* by formula (22)
        Apply **Algorithm 4**, or do the following:
            Recalculate $\{Y_j^m\}_{j=n+1}^N$ for all m using control mappings $\{u_j^*\}_{j=n,\dots,N-1}$
            $V(n, Z_n^m, Y_n^m) = \sum_{j=n}^{N-1} f\big(n, Z_j^m, Y_j^m, u_j^*(Z_j^m, Y_j^m)\big) + g(Z_N^m, Y_N^m) \quad \forall m$
Evaluate the policy

effort of resimulation is proportional to N^2 compared to a linear cost of simulating the process (X_n). Moreover for each step of the resimulated process an optimisation problem has to be solved in order to determine the policy associated to the new trajectory, heavily increasing the computational burden. In this section, we show how this can be avoided.

It follows from (14) in Section 3.2 that in order to avoid resimulation to compute $J(n, (Z_j^m, Y_j^m, u_j^*)_{j=n,\dots,N})$, one needs to be able to reuse $J(n+1, (Z_j^m, Y_j^m, u_j^*)_{j=n+1,\dots,N})$. Notice that an interpolation, analogously as in the grid discretisation approach, it is not a viable alternative because $J(n, (Z_j^m, Y_j^m, u_j^*)_{j=n,\dots,N})$ can not be represented in terms of (X_n, I_n) only as it depends on the whole future trajectory $(Z_j^m, Y_j^m)_{j=n,\dots,N}$.

Here we propose a solution in which we aim to place points Y_n^m in such a way that optimally controlled process reaches the state $(X_{n+1}^m, I_{n+1} = Y_{n+1}^m)$ at time $n+1$ and therefore the existing path $(Z_j^m, Y_j^m)_{j=n+1,\dots,N}$ is a valid successor of (Z_n^m, Y_n^m) enabling reuse of $J(n+1, (Z_j^m, Y_j^m, u_j^*)_{j=n+1,\dots,N})$. Recall that the inventory process follows the dynamics

$$I_{n+1} = \varphi(n, u_n^*(X_n, I_n), I_n).$$

From the optimality condition (22), we obtain that the pathwise optimal antecedent of a point Y_{n+1}^m is a solution $y \in [0, I_{\max}]$, if exists, of the equation

$$Y_{n+1}^m = \varphi\big(n, u_n^*(Z_n^m, y), y\big). \tag{24}$$

In the following lemma, we explore the existence of such y. Define two quantities $\underline{\varphi_n} = \max_{u\in[\underline{u},\overline{u}]} \{\varphi(n, u, 0)\}$ and $\overline{\varphi_n} = \min_{u\in[\underline{u},\overline{u}]} \{\varphi(n, u, I_{\max})\}$ and assume $0 \le \underline{\varphi_n} < \overline{\varphi_n} \le I_{\max}$, $\forall\, n$.

Lemma 1. *A solution $y \in [0, I_{\max}]$ to* (24) *exists if the following conditions are satisfied:*

Algorithm 4 Fixed point performance iteration algorithm

Let y be a solution to $Y_{n+1}^m = \varphi\big(n, u_n(Z_n^m, y), y\big)$.
if solution exists and $y \in [0, I_{\max}]$ **then**
    Set $Y_n^m = y$
    $V(n, Z_n^m, Y_n^m) = f\big(Z_n^m, Y_n^m, u_n^*(Z_n^m, Y_n^m)\big) + V\big(n+1, Z_{n+1}^m, \varphi(n, u_n^*(Z_n^m, Y_n^m), Y_n^m)\big)$
else
    Generate Y_n^m randomly in $[0, I_{\max}]$
    Use the classic resimulation approach or value iteration to compute $V(n, Z_n^m, Y_n^m)$

1. $Y_{n+1}^m \in [\underline{\varphi_n}, \overline{\varphi_n}]$,
2. *the function* $u_n^*(Z_n^m, y)$ *in* (22) *is continuous with respect to* y.

Moreover, if the map $y \mapsto \varphi(n, u_n^*(Z_n^m, y), y)$ *is strictly monotone then the solution is unique.*

Proof. Introduce the function $\psi(y) = Y_{n+1}^m - \varphi(n, u_n^*(Z_n^m, y), y)$ and notice that if condition 2. is satisfied the continuity of ψ is enforced by φ. Evaluate ψ it at the boundaries of $[0, I_{\max}]$:

$$\begin{aligned}
\psi(0) =& Y_{n+1}^m - \varphi(n, u_n^*(Z_n^m, 0), 0) \\
\geq& \underline{\varphi_n} - \underline{\varphi_n} = 0; \\
\psi(I_{\max}) =& Y_{n+1}^m - \varphi(n, u_n^*(Z_n^m, I_{\max}), I_{\max}) \\
\leq& \overline{\varphi_n} - \overline{\varphi_n} = 0.
\end{aligned}$$

Therefore, if condition 1. is satisfied it must exists a point $y \in [0, I_{\max}]$ such that $\psi(y) = 0$, i.e., solve equation (24). In addition, if $y \mapsto \varphi(n, u_n^*(Z_n^m, y), y)$ is strictly monotone then so is ψ and hence there exists only one y which solves equation (24). □

Lemma 1 ensures the existence of solution to the equation (24) even though such solution could be found also for weaker assumptions, possibly outside the interval $[0, I_{\max}]$. Whenever this solution does not lie in $[0, I_{\max}]$, i.e. $Y_{n+1}^m \notin [\underline{\varphi_n}, \overline{\varphi_n}]$, it is not possible to determine a step backward Y_n^m for the inventory variable. In such case, we propose to break the path and position the point Y_n^m at random in $[0, I_{\max}]$. Following this step, we either

- resimulate the path $\{Y_j^m\}_{j=n+1,\dots N}$ as in the traditional performance iteration algorithm, or
- replace $J(n+1, (Z_j^m, Y_j^m, u_j^m)_{j=n+1,\dots,N})$ by an approximation of the value function

$$\sum_{k=1}^{K} \alpha_k^{n+1} \hat{\phi}_k\big(Z_{n+1}^m, \varphi(n, u_n^*(Z_n^m, Y_n^m), Y_n^m)\big)$$

  analogously as in the value iteration algorithm.

The random positioning of training points Y_n^m when the solution of the constrained fixed point problem (24) cannot be found avoids the clustering of training points in certain areas of the state space.

The phenomenon described above that there is no solution to (24) that satisfies the bounds on the inventory level is not unusual in many practical examples and, in particular, does not result from approximation errors or any other features of the numerical scheme. In an example in the next section, we consider trading a commodity, with price (X_t), subject to constraints on the size of storage. An

optimal behaviour when the price is high is to sell. However, if this happens, an optimal trading would never lead to a full inventory at the next time period, $Y^m_{n+1} = I_{\max}$. In fact, with high commodity prices over a period of time and an optimal strategy of selling, the backward constructed trajectory Y^m_j would climb up (when time goes backwards) finally hitting the level $I_{\max}$.

Example. It is commonly seen in the inventory problems that the control represents additive adjustments to the inventory level, i.e. $\varphi(n, u, i) = i + u$. Equation (24) simplifies to

$$Y^m_{n+1} = y + u^*_n(Z^m_n, y).$$

Denoting $\Delta y = Y^m_{n+1} - y$ we obtain a classical fixed point problem

$$\Delta y = \hat{\nu}(\Delta y),$$

where $\hat{\nu}(\Delta y) = u^*_n(Z^m_n, Y^m_{n+1} - \Delta y)$. Lemma 1 implies that exists an unique solution $\Delta y \in [\underline{u}, \overline{u}]$ for every $Y^m_{n+1} \in [\overline{u}, I_{\max} + \underline{u}]$.

In the following section we evaluate our methods on two examples and compare them against their main competitors: grid discretisation and control randomisation.

5 Numerical evaluation

This section collects numerical results showcasing the performance of our regress-later approach. Our first example is a "price arbitrage" problem in the context of energy trading for which we compare the accuracy and running time of performance iteration algorithm using grid discretisation (GD), control randomisation (CRP) and the regress-later Monte Carlo approach. We consider three variations of the latter: resimulation of inventory after each backward step (RLMCP) and backward construction of paths for performance iteration with the resimulation of broken paths (RLMC BP) and with the use of value function approximation (RLMC BV).

Our second control problem is multi-dimensional and consists in controlling the operations of a power system comprised of a wind turbine, a group of residential buildings, an energy storage and an interconnection with the main grid with the aim of maximising the profit over a finite horizon. Due to the dimensionality the underlying process and a multi-dimensional control, we apply regress-later algorithm only. The multi-dimensionality of the control prevents usage of grid discretisation while making control randomisation significantly slower than RLMC.

5.1 Price arbitrage problem

Consider the following specification of the dynamics of a commodity price as an Ornstein-Uhlenbeck process (mean-reverting to 0 for simplicity) and inventory in continuous time (with time counted in years):

$$\begin{cases} dX_t = -10X_t\,dt + 3\,dW_t, \\ dI_t = u_t\,dt\,, \;\; u \in \mathcal{U}, \end{cases} \tag{25}$$

where $\mathcal{U} = \{u_s \in \{-10, 0, 10\} : \; 0 \le I_t \le 1 \;\forall t\}$ consists of adapted processes. The objective is to maximise the functional

$$\mathbb{E}\Big[\int_t^T -(u_s + \rho_{u_s})\,X_s\,ds + \frac{1}{2}X_N\,I_N | X_t = x, I_t = i\Big], \tag{26}$$

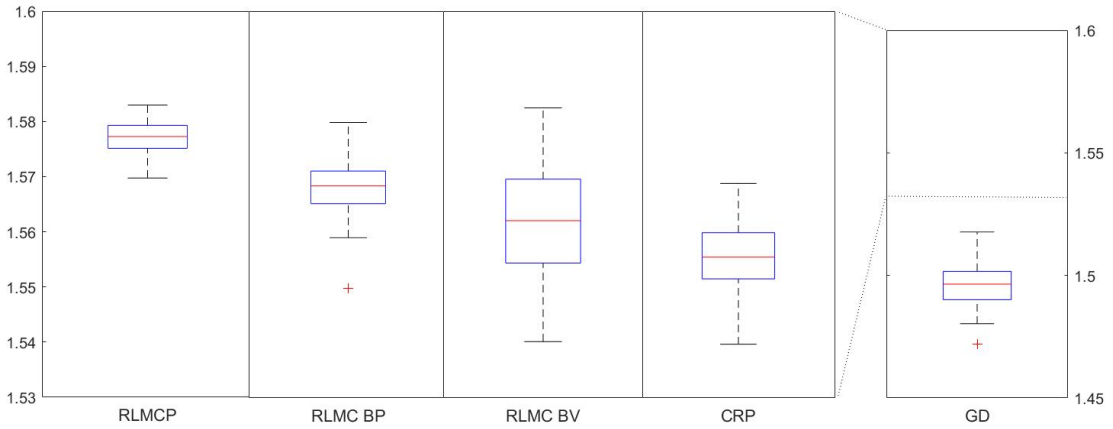

Figure 1: Estimators of the value $V(0,0,0.5)$ of the policies computed by: Regress Later Monte Carlo Performance iteration (RLMCP), with Backward construction of paths (RLMC BP) and using Value iteration approximations for broken paths (RLMC BV); Control Randomisation Performance iteration (CRP); Grid Discretisation of the inventory (GD). Each boxplot is obtained by repeating 50 times the policy estimation with $M = 1000$ while the Monte Carlo evaluation of policites is computed over the same set of $M' = 100,000$ trajectories. The 95% confidence interval for all policies is narrower than ± 0.004.

where ρ_{u_s} introduces a cost associated with each control: $\rho_{u_s} = 1_{\{u_s \neq 0\}} 2$.

We discretise this problem as in Section 2.1 with $T = 1$, $\Delta = 1/100$, and consequently $N = 100$. The resulting dynamic programming equation is of the form

$$V(t,x,i) = \max_{\substack{u \in \{10,0,-10\}, \\ i+u\Delta \in [0,1]}} \Big\{ -(u+\rho_u)\, X_n \Delta + \mathbb{E}\big[V(n+1, X_{n+1}, i+u) | X_n = x \big] \Big\}.$$

5.1.1 Implementation

All algorithms are implemented in MatLab and are optimised in a similar fashion in order to reveal the relative differences in their computational complexity. In particular, we do not take advantage of vectorisation whenever this speed-up cannot be shared by all algorithms. The basis in RLMC algorithms consisted of polynomials up to order two for X and order three for I, i.e. $\{1,\, x,\, i,\, xi,\, x^2,\, i^2,\, x^2 i,\, x^2 i^2,\, i^3\}$. Basis functions for control randomisation involve also the control variable, so in order to use the same number of basis function we employed a basis comprising polynomials up to order two for X, I and u, i.e. $\{1,\, x,\, i,\, u,\, xi,\, x^2,\, i^2,\, u^2,\, xu,\, iu\}$. For the GD method we discretised the inventory in $L = 10$ levels and then used polynomial basis functions up to order two for X, i.e. $\{1,\, x,\, x^2\}$. Despite this small basis, the implementation of GD algorithm was still significantly slower than all the alternative methods.

5.1.2 Comparison

We tested the accuracy of the methods running 50 separate computations of $V(0,0,0.5)$ choosing $M = 1000$ and evaluating the resulting policies over a the same set of $100,000$ simulations shared between all methods and runs in order to minimise the effect of the Monte Carlo error on relative results. The 95% confidence intervals for the reported performance of computed policies are below

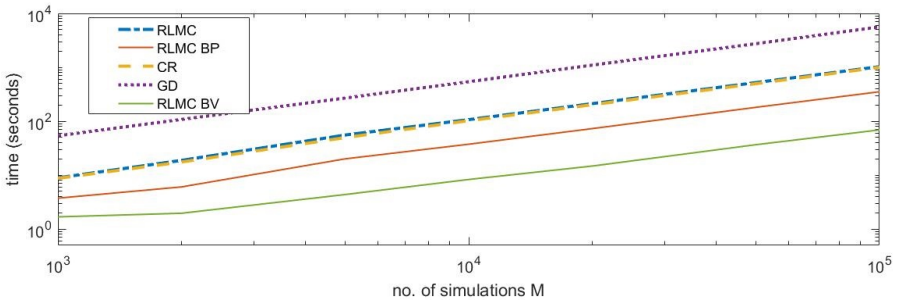

Figure 2: Computational time for policy estimation as function of M. Both scales are logarithmic and time is counted in seconds.

± 0.004. Boxplots of objective values are shown in Figure 1. A tall box indicates that policies computed in different runs may give significantly different expected profits/losses. Notice that RLMCP has superior accuracy to all other methods. Its improved performance over RLMC BP may be explained by resimulation present in the former while the latter uses the same training points for both estimation of the optimal policy and its evaluation; this drawback is well known in literature. Loss of accuracy when backward inventory construction uses estimates of value function on broken paths (RLMC BV) instead of resimulation is due to the replacement of true evaluation of policy with an estimate. Control randomisation (CRP) yielded lower objective value and larger variation compared to RLCMP which uses the same method of path resimulation; these results were also inferior to those produced by other RLMC methods. Finally, GD performs significantly worse than all other methods. We believe this poor performance is due to inconsistencies in the estimation error on the regression coefficients for different layers of inventory.

The running time of the algorithms as a function of M is shown in Figure 2. It can be seen that the computational complexity is linear in M with different methods having different proportionality factors reflected in the parallel shift of the lines on the graph. The use of backward construction of the paths improve, for this particular problem, the constant to about 25% of the original value as approximately one quarter of paths are broken at each time step. If we choose the value iteration update for those broken paths achieve significantly better running time. Note that even though GD uses a direct update for the value function it is the slowest method because of the multiple regressions computed at each time step. It is easy to deduce from the algorithms that the computational complexity with respect to the number of time periods N is quadratic for methods with path resimulation (CRP, RLMCP, RLMC BP) and linear for others (RLMC BV, GD).

To conclude, in order to give a feeling of the difference between the algorithms we propose a comparison, in Figure 3, of a particular path of the inventory and the associated controls obtained with each method. This comparison is not exhaustive as it only shows one instance of random behaviour but it may suggest to what extent the estimated policies can differ. The bottom graph displays controls at each time step with colours corresponding to an algorithm used for the calculation of the policy and indicated in the legend in the middle graph. The controls are *buy, store, sell* for $u = 0.1, 0, -0.1$. Poor performance of grid discretisation algorithm can be understood by observing, for example, times 50-60. At times 50-55 all policies agree to store because the price is high and the inventory is empty. However, at times 56-57 the price drops and all methods apart from GD switch to buying.

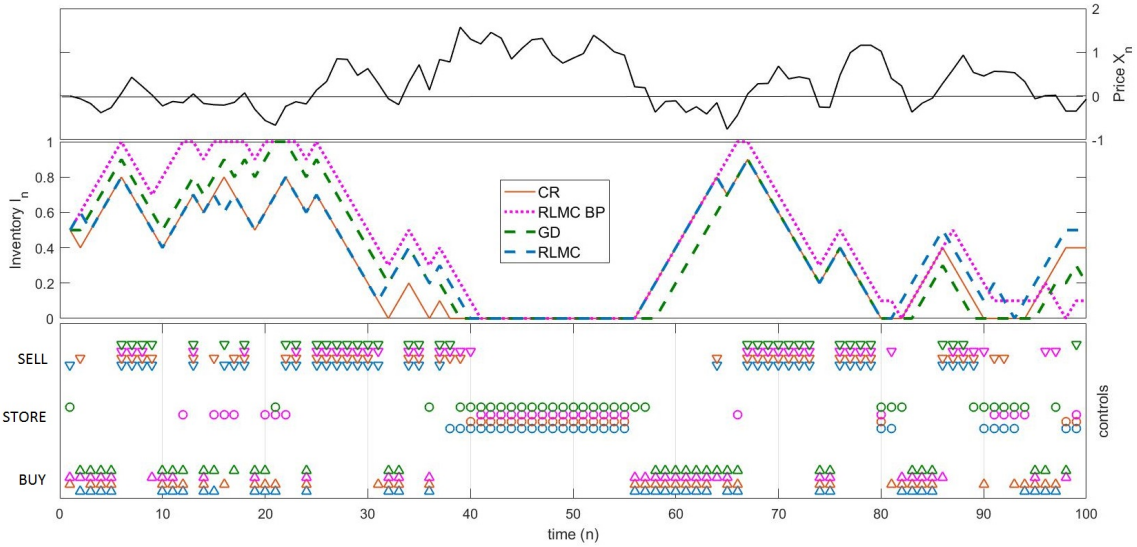

Figure 3: The graph above displays (from the top) a realisation of the process X, the trajectories of the inventory and the control used at each time step. This representation can suggest to what extent the different policies produce different behaviours.

5.2 Energy Storage management

In this subsection we present a high dimensional problem in which the control is initially five-dimensional and the state space four-dimensional; we will reduce the action space to two dimensions and solve the problem applying $RLMC$ algorithm.

For the model in this example we take inspiration from Scott et al. (2014) and consider an electricity system composed of an intermittent energy supply represented by a wind turbine, a load or electricity demand represented by a group of residential buildings, a battery storage device and an interconnection with the main grid. Energy can flow through the different components of the system as indicated by the arrows in Figure 4; note that the power produced from the wind turbine is directly transferred to the load and used to satisfy the demand, the power in excess is transferred to the storage device which can sell it to the main grid or save it for later use in the battery. When the power from the wind turbine is insufficient to satisfy the demand, the shortfall is fulfilled either (or both) from the storage device discharging the battery or/and by buying electricity from the grid. Note that we assume that the main grid can always provide the power we require.

The operator has control on every flow of energy, making the dimensionality of the control process equal to five; we denote these flows at time n by X_n^{WD},X_n^{WB},X_n^{BD},X_n^{GB} and X_n^{GD}, where the notation $X^{i,j}$ indicates energy flowing from i to j and we denote the components as follow: W for the wind turbine, D for the demand, G for the grid interconnection and B for the battery. We constrain each flow to be positive except for X_n^{GB} which can assume negative values indicating that the battery is selling energy to the grid. Finally we consider our problem stated on an hourly basis over a time horizon of two weeks.

As in Scott et al. (2014) we make the following modelling choices.

Wind Power

We compute the power in MWh produced by a wind turbine with rated power $2.1\,MW$ in one hour as a function of the wind speed w_n (average over an hour in m/s):

$$W_n = 10^{-6}\, C_p\, \rho\, A\, \max(w_n, \omega)^3,$$

where $C_p = 0.4$ is the power efficiency, $\rho = 1.225\, kg/m^3$ is the density of air, $\omega = 14m/s$ represents the rated output speed and $A = 50^2\pi$ is the area swept by the blades.

In particular we model the centred square root of the wind speed (which we consider constant) as an AR(1) process:

$$\begin{aligned} w_n &= (y_n + \mu)^2 \\ y_{n+1} &= 0.7633 y_n + 0.4020\xi_n, \quad \xi_n \sim \mathcal{N}(0,1) \end{aligned}$$

The coefficients are calibrated from 2010 real data registered in Maryneal, Texas; see Scott et al. (2014) for details.

Electricity Price

The electricity price is modelled as is Scott et al. (2014) as an exponential of a mean reverting process superimposed on a weekly seasonal component Y^{week} estimated from the data recorded in Texas in 2010:

$$\begin{aligned} P_{n+1} &= e^{Y_{n+1}} - 27.2531 + Y^{week}_{n+1} \\ Y_{n+1} &= Y_n + 0.2055(4.1995 - Y_n) + 0.11856\xi_n + J_n, \quad \xi_n \sim \mathcal{N}(0,1) \end{aligned}$$

where $J_n = \xi^j_n 1_{\{U_n < 0.017\}}$ is the jump component, $\xi^j_n \sim \mathcal{N}(0, 0.4229)$ and $U_n \sim \mathcal{U}(0,1)$.

Note that we introduce a fixed 10% transmission surcharge on the use of the grid for both supplying demand and exchanging energy with the battery. This is a requirement in order to make the optimal control unique, otherwise it is equivalent to supply power to the load from the battery through X^{BD} or through $X^{GB} \to X^{GD}$.

Electricity Demand

It is well known that temperature is one of the main drivers of the electricity demand of residential buildings, in fact during periods of high or low temperature households tend to use air conditioning. The electricity demand is therefore modelled as a function of temperature; in Scott et al. (2014) it is suggested to use an order 5 polynomial to fit the observed pairs of demand and temperature. Since estimated coefficients were not provided and in order to be consistent with the other data, we estimated our coefficients using observations recorded in Austin, Texas in 2010 (`http://www.noaa.gov/`). We obtained the following model for the demand:

$$\tilde{D}_n = 6784.9728 - 235.5911\, T_n - 2.2869\, T_n^2 + 0.8897\, T_n^3 - 0.0204\, T_n^4 + 0.000105\, T_n^5.$$

To scale down the demand to a few buildings we set $D_n = \frac{\tilde{D}_n}{30000}$. We model the temperature as a simple mean reverting process $T_n = \tilde{T}_n + \mu_n$, where μ_n represents a daily seasonal component estimated from the data and

$$\tilde{T}_{n+1} = -0.92\tilde{T}_n + 2.14\xi_n, \quad \xi_n \sim \mathcal{N}(0,1).$$

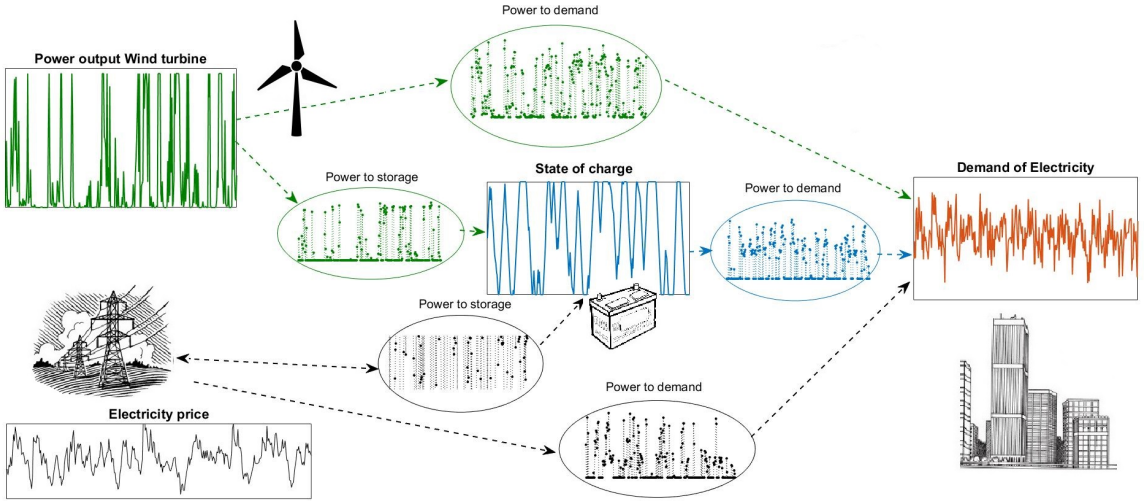

Figure 4: The diagram above shows components of the system. Boxes show evolution of state variables. Elipses display energy flows among the components. Note that the algorithm directly controls the connection grid-battery and grid-demand as other flows are uniquely determined by those two.

The Battery

The most interesting component of the system is an energy storage device; it is represented by a battery with rated capacity $I_{\max} = 20$MWh and maximum power in charge $\Pi_{\max} = 2.1$MW and discharge $\Pi_{\min} = -2.1$MW. For simplicity of exposition we assume the battery is 100% efficient and its use does not take into account its ageing. The level of charge of the battery evolves as follow:

$$I_{n+1} = I_n + X_n^{GB} + X_n^{WB} - X_n^{BD}.$$

5.2.1 Problem formulation and solution

We look at the system on an hourly basis during the first two weeks of January, which corresponds to 336 hourly time steps. We aim to provide power to the residential buildings minimising the cost of buying electricity from the grid. Since the problem is stated in discrete time we will deal with energy rather than power. Notice that the following set of constraints is in place (we dropped the subscript n):

$$X^{GD} + X^{WD} + X^{BD} = D \tag{c.1}$$
$$X^{WD} = \min\{W, D\} \tag{c.2}$$
$$X^{WB} + X^{WD} = W \tag{c.3}$$
$$X^{GD} \geq 0 \tag{c.5}$$
$$X^{BD} \in [0, -\Pi_{\min}] \tag{c.6}$$
$$X^{GB} + X^{WB} - X^{BD} \in [\Pi_{\min}, \Pi_{\max}] \cap [-I, I_{\max} - I] \tag{c.7}$$

The first constraint (c.1) is natural and says that the demand has to be matched at all times. Constraints (c.2) and (c.3) regulate the operations of the wind turbine: all its power is delivered to the load first and the excess is transferred to the battery. As long as the current level of charge allows it the

power is accumulated; the excess power is sold to the grid (for a price reduced by the 10% transmission surcharge), c.f. (c.7). Constraint (c.5) enforces the natural condition that power should only flow from the grid to the load. Constraint (c.6) limits the discharging rate of the battery enforcing a bound on the energy flowing from the battery to the demand. Finally the last constraint (c.7) guarantees that the inventory stays between empty and full and that the maximum charging/discharging rates are maintained.

Note that the constraints (c.1)-(c.7) actually reduce the dimension of the control process from five to two (X^{GB} and X^{GD}), as the other components can be computed consequently. Recall that $X^{GD} \geq 0$ while the sign of X^{GB} depends on whether the battery buys or sells energy to the grid.

We can then write the profit of an electricity provider who owns the wind turbine and the battery and has contracted a given load, as electricity sold minus electricity bought (net of transmission surcharge):

$$J\left(\left(D_s, W_s, P_s, B_s, I_s, X_s^{GB}, X_s^{GD}\right)_{s=1,\dots,N}\right) = \sum_{n=1}^{N} P_n\, D_n - 1.1 P_n (X_n^{GB} + X_n^{GD}).$$

Define the value function of the problem as the solution to the minimisation problem:

$$V(n,d,w,p,i) = \min_{X \in \mathcal{U}} \left\{ -\mathbb{E}\Big[J\big((D_s, W_s, P_s, B_s, I_s, X_s^{GB}, X_s^{GD})_{s=n,\dots,N}\big) | (D_n,\, W_n,\, P_n,\, I_n) = (d,w,p,i)\Big]\right\}, \tag{27}$$

where $\mathcal{U}$ is the set of admissible controls specified by constraints (c.1)-(c.7).

5.2.2 Numerical solution

We solve the problem (27) applying the Regress Later algorithm introduced in Section 4. For this problem the choice of linear and quadratic basis functions, i.e. $\{D,\, W,\, P,\, D^2,\, W^2,\, P^2,\, I,\, D\,I,\, P\,I,\, W\,I,\, I^2\}$, works satisfactorily. We could also not see any significant improvement in results when using more than 5000 simulated trajectories. The numerical complexity of the algorithm stems mainly from a non-trivial two dimensional constrained optimisation problem for the control which is solved at each time step and for each training point.

Once the estimated optimal policy has been computed we run it forward in time to assess its value; in Figure 4 we display a particular realisation of all the flows of energy among the components. It can be observed that the battery stores the wind power when this exceeds the demand of electricity and transfers it to the demand instead of buying electricity from the grid at high price. Over the two weeks of interest the system considered has an average demand of electricity of 320 KW compared with an average production of wind power of 350 KW meaning that the system is a net exporter of electricity. Figure 6 displays some of the interactions within the system. The scatter plot on the left shows the energy flowing to the demand from the battery (blue) and from the grid (orange). There is a general tendency for the power to be bought from the grid when its price is low while it is supplied mainly by the battery when the price is high. The scatter plot on the right displays the interactions between wind power and transfers through the connection battery-grid. We observe that when the wind power is high the battery tends to sell energy to the grid in order to maintain an acceptable level of inventory.

The solution of the optimisation problem of this subsection enables study of the economic value of using a battery to assist the operations of a wind turbine. To this end, we evaluate the estimated policy on a new and larger set of trajectories of the underlying processes (wind, demand and price). This is

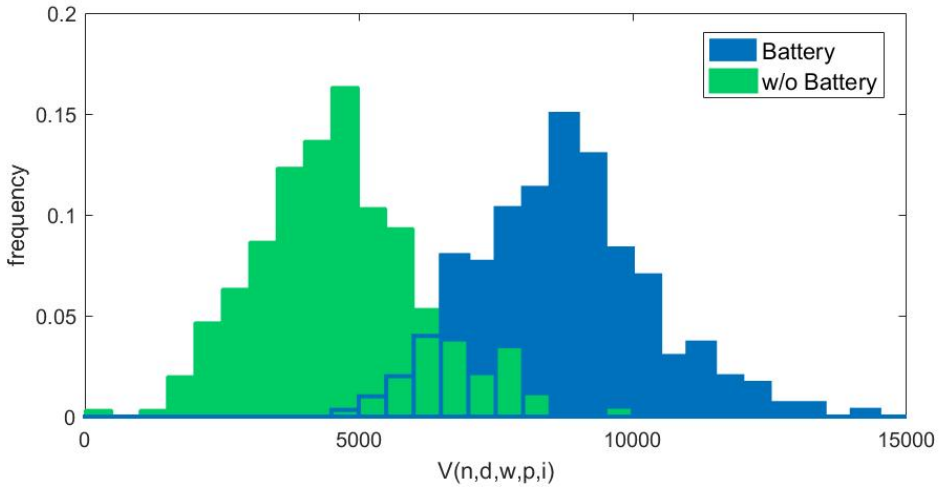

Figure 5: The empirical probability distributions of the performance J at time 0 with the initial data $(d, w, p, i) = (0.15, 0, 35, 10)$, representing the profits obtained in the system with (blue) and without (green) the use of the battery. There are 30 bins each 500 units wide.

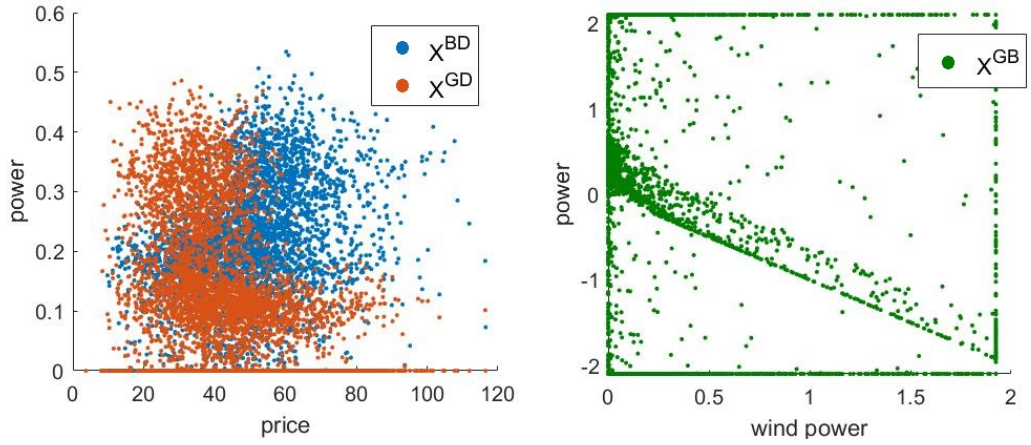

Figure 6: This graphs display interactions within the system in the form of scatter plots. The left plot shows X^{GD} and X^{BD} against the price revealing that the first tends to provide power when price is low, the latter when price is high. The right plots show how the connection battery-grid (X^{GB}) is influenced by the wind power; note that the high level of wind often results in selling energy to the grid in order to decrease the level of charge.

further compared to the profits made in a system with no battery installed, i.e., all excess wind energy immediately sold to the grid. We find that the 25% and 75% quantiles of the profits for operating the system without the use of the storage over two weeks are 3595, and 5498, respectively. When the flexibility provided by the storage is optimally used the respective quantiles are 7556 and 9598, showing an increase in the profits of around 4000. Further details about the distribution of profits are displayed in Figure 5. Considering the inventory behaviour displayed in Figure 4 as "typical" we can compute the total throughput of energy of 490 MWh in the two weeks. Assuming further that the battery sustains 4000 full cycles we obtain an estimated life-span of 320 weeks or approximately 6 years. Assuming further a stationary flow of profit during the life of the battery we conclude that the installation of the storage alongside the wind turbine will bring additionally $\$640000$ in profits (assuming zero scrap value and zero interest rate).

Our solution of problem (27) shows that RLMC is well suited for dealing with the optimal control of complex systems and therefore provides an alternative to heuristics and policies obtained under perfect foresight used commonly for the solution of real world high dimensional problems.

6 Conclusions

We presented Monte Carlo methods for the optimal control of an inventory subject to a measure of performance based on an exogenous random process. Our main contribution was in designing numerical optimisation algorithms based on regress later approach and using value and performance iteration schemes. In particular, we showed how, solving a fix point problem, one can avoid resimulation of the inventory trajectories in performance iteration during the backward recursion steps. Numerical tests demonstrated that our regress-later Monte Carlo methods perform well compared to competitors and are suitable for solving real-world high-dimensional problems.

References

Balata, A. and Palczewski, J. (2017). Regress-later Monte Carlo for controlled Markov chains. in preparation.

Beutner, E., Pelsser, A., and Schweizer, J. (2013). Fast convergence of regress-later estimates in least squares Monte Carlo. *SSRN 2328709*.

Boogert, A. and de Jong, C. (2008). Gas storage valuation using a Monte Carlo method. *The Journal of Derivatives*, 15(3):81–98.

Bouchard, B. and Warin, X. (2012). Monte-Carlo valuation of american options: Facts and new algorithms to improve existing methods. In Carmona, R. A., Moral, P. D., Hu, P., and Oudjane, N., editors, *Numerical Methods in Finance*, pages 215–255. Springer.

Broadie, M. and Glasserman, P. (2004). A stochastic mesh method for pricing high-dimensional american options. *Journal of Computational Finance*, 7(35):35–72.

Broadie, M., Glasserman, P., and Ha, Z. (2000). Pricing American options by simulation using a stochastic mesh with optimized weights. In Uryasev, S. P., editor, *Probabilistic Constrained Optimization*, volume 49 of *Nonconvex Optimization and Its Applications*, pages 26–44. Springer.

Carmona, R. and Ludkovski, M. (2010). Valuation of energy storage: an optimal switching approach. *Quantitative Finance*, 10(4):359–374.

Carriere, J. F. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics*, 19(1):19–30.

Delage, E., Denault, M., and Simonato, J.-G. (2017). A simulation-and-regression approach for dynamic programming, and its application to portfolio choice. *Optimization and Engineering*, forthcoming.

Denault, M., Simonato, J.-G., and Stentoft, L. (2013). A simulation-and-regression approach for stochastic dynamic programs with endogenous state variables. *Computers & Operations Research*, 40(11):2760–2769.

Glasserman, P. and Yu, B. (2002). Simulation for american options: regression now or regression later? In Niederreiter, H., editor, *Monte Carlo and Quasi-Monte Carlo Methods*, pages 213–226. Springer Berlin Heidelberg.

Guthrie, G. A. (2009). *Real options in theory and practice*. Oxford University Press.

Haessig, P., Kovaltchouk, T., Multon, B., Ahmed, H. B., and Lascaud, S. (2013). Computing an optimal control policy for an energy storage. *arXiv:1404.6389*.

Jain, S. and Oosterlee, C. W. (2012). Pricing high-dimensional Bermudan options using the stochastic grid method. *International Journal of Computer Mathematics*, 89(9):1186–1211.

Jain, S. and Oosterlee, C. W. (2015). The stochastic grid bundling method: Efficient pricing of bermudan options and their greeks. *Applied Mathematics and Computation*, 269:412–431.

Kharroubi, I., Langrené, N., and Pham, H. (2014). A numerical algorithm for fully nonlinear HJB equations: an approach by control randomization. *Monte Carlo Methods and Applications*, 20(2):145–165.

Kitapbayev, Y., Moriarty, J., and Mancarella, P. (2015). Stochastic control and real options valuation of thermal storage-enabled demand response from flexible district energy systems. *Applied Energy*, 137:823–831.

Kushner, H. J. and Dupuis, P. G. (2001). *Numerical Methods for Stochastic Control Problems in Continuous Time*. Applications of Mathematics. Springer-Verlag New York.

Langrené, N., Tarnopolskaya, T., Chen, W., Zhu, Z., and Cooksey, M. (2015). New regression Monte Carlo methods for high-dimensional real options problems in minerals industry. 21st International Congress on Modelling and Simulation, Gold Coast, Australia, 29 Nov to 4 Dec 2015 www.mssanz.org.au/modsim2015.

Longstaff, F. A. and Schwartz, E. S. (2001). Valuing american options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1):113–147.

Nadarajah, S., Margot, F., and Secomandi, N. (2017). Comparison of least squares Monte Carlo methods with applications to energy real options. *European Journal of Operational Research*, 256:196–204.

Scott, W. R., Powell, W. B., and Moazeni, S. (2014). Least squares policy iteration with instrumental variables vs. direct policy search: Comparison against optimal benchmarks using energy storage. *arXiv:1401.0843*.

Secomandi, N. (2010). Optimal commodity trading with a capacitated storage asset. *Management Science*, 56(3):449–467.

Swindle, G. (2014). *Valuation and Risk Management in Energy Markets*. Cambridge University Press.

Tsitsiklis, J. N. and VanRoy, B. (2001). Regression methods for pricing complex american-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703.

Zhaoa, G. and Davison, M. (2009). Optimal control of hydroelectric facility incorporating pump storage. *Renewable Energy*, 34(4):1064–1077.